

BI-LIPSCHITZ EXTENSION FROM BOUNDARIES OF CERTAIN HYPERBOLIC SPACES

ANTON LUKYANENKO

ABSTRACT. Tukia and Väisälä showed that every quasi-conformal map of \mathbb{R}^n extends to a quasi-conformal self-map of \mathbb{R}^{n+1} . The restriction of the extended map to the upper half-space $\mathbb{R}^n \times \mathbb{R}^+$ is, in fact, bi-Lipschitz with respect to the hyperbolic metric.

More generally, every homogeneous negatively curved manifold decomposes as $M = N \rtimes \mathbb{R}^+$ where N is a nilpotent group with a metric on which \mathbb{R}^+ acts by dilations. We show that under some assumptions on N , every quasi-symmetry of N extends to a bi-Lipschitz map of M .

The result applies to a wide class of manifolds M including non-compact rank one symmetric spaces and certain manifolds that do not admit co-compact group actions. Although M must be Gromov hyperbolic, its curvature need not be strictly negative.

1. INTRODUCTION

A classical result of Heintze states that any homogeneous manifold M with negative sectional curvature is diffeomorphic to a Lie group $N \rtimes \mathbb{R}^+$. Here N is a nilpotent group diffeomorphic to \mathbb{R}^n for some n , and \mathbb{R}^+ acts on N by a family of diffeomorphisms α_s . As a negatively curved space, M possesses a boundary at infinity $\partial_\infty M$, composed of geodesic rays up to appropriate equivalence. The abstract boundary $\partial_\infty M$ may be identified with the topological boundary of M as a subset of $(N \times \mathbb{R}) \cup \{\infty\}$; namely, $\partial_\infty M \cong N \cup \{\infty\}$. This identification provides N with a family of parabolic visual metrics; for one of the visual metrics, the diffeomorphism α_s is a homothety with dilation factor s .

A map $f : M \rightarrow M$ is a quasi-isometry if there exist $L \geq 1, C \geq 0$ such that for all $x, y \in M$:

$$(1.1) \quad -C + L^{-1}d_M(x, y) \leq d_M(fx, fy) \leq Ld_M(x, y) + C.$$

It is known that every quasi-isometry $f : M \rightarrow M$ extends to a quasi-symmetry $\partial_\infty f : \partial_\infty M \rightarrow \partial_\infty M$. Conversely, every quasi-symmetry of $\partial_\infty M$ extends to a quasi-isometry of M , unique up to bounded additive error.

A standard problem is to modify f , up to bounded error, so that in (1.1) one has $C = 0$ or $L = 1$. In the first case, f is replaced by a bi-Lipschitz map, and in the second by a rough isometry. In very specialized cases such as Mostow's rigidity theorem, one accomplishes both goals simultaneously: f is bounded distance from an isometry.

Date: December 13, 2011.

2010 Mathematics Subject Classification. 30L10, 53C23, 53C30.

In their 1982 paper, Tukia and Väisälä examined the bi-Lipschitz approximation problem for real hyperbolic space $\mathbb{H}_{\mathbb{R}}^n$.

Theorem 1.1 (Tukia-Väisälä [24]). *Let $f : \mathbb{H}_{\mathbb{R}}^{n+1} \rightarrow \mathbb{H}_{\mathbb{R}}^{n+1}$ be a quasi-isometry. Then f is bounded distance from a bi-Lipschitz map $F : \mathbb{H}_{\mathbb{R}}^{n+1} \rightarrow \mathbb{H}_{\mathbb{R}}^{n+1}$.*

Remark 1.2. The main theorem of [24] is, in fact, that every quasi-symmetry g of \mathbb{R}^n extends to a quasi-symmetry G of \mathbb{R}^{n+1} . This follows from Theorem 1.1 by extending g to a quasi-isometry of $\mathbb{H}_{\mathbb{R}}^{n+1}$ and approximating it by a bi-Lipschitz map. One then observes that, in the upper half-space model, $\mathbb{H}_{\mathbb{R}}^{n+1}$ is conformally equivalent to $\mathbb{R}^n \times \mathbb{R}^+$ to complete the proof. See Section 3.1 for more details.

Our goal is to generalize Theorem 1.1 to a larger class of hyperbolic spaces. Let M be a Gromov hyperbolic manifold. The following problems are equivalent:

- (1) Is every quasi-isometry of M bounded distance from a bi-Lipschitz map?
- (2) Does every quasi-symmetry of $\partial_{\infty}M$ extend to a bi-Lipschitz map of M ?

Following Tukia-Väisälä, we focus on version (2) of the problem.

The dyadic tiling of the upper half-space (see Section 3.1) is the essential combinatorial ingredient in the proof of Theorem 1.1. The tiles are related to each other by isometries of hyperbolic space, and the intersection graph of the tiling captures the geometry of $\mathbb{H}_{\mathbb{R}}^{n+1}$. We define two classes of spaces (see Section 3.2 for precise definitions) that possess properties of \mathbb{R}^n and $\mathbb{H}_{\mathbb{R}}^{n+1}$ essential to the proof of Theorem 1.1.

We say that a metric space H homeomorphic to \mathbb{R}^n possesses a *stacked tiling* if it has an analogue of the integer lattice. Namely, H must admit a discrete cocompact group of isometries Γ and a homothety α . Furthermore, we require that there exists a fundamental domain K for the action of Γ whose rescaling αK is tiled by translates of K . The space $H \times \mathbb{R}^+$, endowed with a Riemannian metric, is a *metric similarity space* if both Γ and α extend appropriately to isometries of H^+ . We say that H is the *base* of H^+ .

One can extend the stacked tiling of the base H to a tiling of H^+ by dyadic-type tiles. We show that H^+ is Gromov hyperbolic and define the *vertical direction* as the point $[\gamma_{\infty}] \in \partial_{\infty}H^+$ represented by quasi-geodesics of the form $\{x\} \times [1, \infty)$, for $x \in H$. We then identify $\partial_{\infty}H^+ \setminus \{[\gamma_{\infty}]\}$ with $H \times \{0\}$ as sets, and show that the metric d_H on H is, under the identification, a parabolic visual metric on $\partial_{\infty}H^+ \setminus \{[\gamma_{\infty}]\}$. We conclude that any quasi-isometry of H^+ preserving the vertical direction extends to a quasi-symmetry of H .

We then prove (see Section 3.7 for the precise statement):

Theorem 1.3. *Let $H^+ = H \times \mathbb{R}^+$ be a metric similarity space of dimension not equal to 4. Then every quasi-symmetry of H extends to a bi-Lipschitz mapping of H^+ . Equivalently, every quasi-isometry of H^+ preserving the vertical direction is bounded distance from a bi-Lipschitz map.*

Theorem 1.3 applies, in particular, to the homogeneous manifolds $M = N \rtimes \mathbb{R}^+$ where the group structure of the nilpotent group N is defined by polynomials with rational coefficients and the Lie algebra of N is graded (see Definition 4.3).

Theorem 1.4. *Let N be a nilpotent group with rational coefficients and graded Lie algebra, with the corresponding Carnot-Carathéodory metric. Letting \mathbb{R}^+ act on M by homotheties, set $M = N \rtimes \mathbb{R}^+$ and give M a left-invariant Riemannian metric. If the dimension of M is not equal to 4, then any quasi-isometry of $N \rtimes \mathbb{R}^+$ preserving the vertical direction is bounded distance from a bi-Lipschitz map.*

Remark 1.5. Examples of nilpotent groups with rational coefficients and graded Lie algebras include the Heisenberg group and higher-order jet spaces. We may take M to be any non-compact rank one symmetric space of dimension not equal to 4, including complex hyperbolic spaces $\mathbb{H}_{\mathbb{C}}^n$, $n \neq 2$. Note that for $\mathbb{H}_{\mathbb{C}}^n$, the condition of preserving the vertical direction is not necessary because the isometry group acts transitively on $\partial_{\infty} \mathbb{H}_{\mathbb{C}}^n$.

Another generalization of Theorem 3.1 was given recently by X. Xie.

Theorem 1.6 (Xie [25]). *Let \widetilde{M} be the universal cover of a compact negatively curved manifold of dimension not equal to 4. Then every quasi-isometry of \widetilde{M} is bounded distance from a bi-Lipschitz map.*

Both Theorem 1.3 and Theorem 1.6 apply to spaces with rich isometry groups, such as $\mathbb{H}_{\mathbb{C}}^n$. Apart from the core components, both the assumptions and proofs of the two theorems are quite different, as we discuss in Section 4.3.

The structure of the paper is as follows. We first review the theory of Gromov hyperbolic spaces and their boundaries in Section 2. In Section 3 we summarize the proof of Theorem 3.1, define and study metric similarity spaces, and prove Theorem 1.3. In Section 4, we provide examples of metric similarity spaces and compare our result to Theorem 1.6, concluding with some remarks and open questions.

I would like to thank my PhD adviser Jeremy Tyson for his help and guidance on this project, as well as Richard Schwartz and John Mackay for interesting conversations on the subject. The paper also benefited highly from the 2011 summer school on filling invariants and asymptotic structures organized by David Fisher.

2. GROMOV HYPERBOLIC GEOMETRY

2.1. Basics. We briefly review the theory of Gromov hyperbolic spaces. For a more complete exposition, see [5, 2, 10] and references therein.

For $C > 0$, a C -rough isometry between two metric spaces (X, d_X) and (Y, d_Y) is a map $f : X \rightarrow Y$ such that

$$(2.1) \quad -C + d_X(x, y) \leq d_Y(fx, fy) \leq d_X(x, y) + C$$

for all $x, y \in X$. Note that a rough isometry need not be continuous.

A $(C$ -rough) geodesic in a metric space X is an $(C$ -rough) isometry $f : I \rightarrow X$, where I is an interval. The space X is $(C$ -roughly) geodesic if any two points are joined by a $(C$ -rough) geodesic. If X is C -roughly geodesic and C is clear from context, we denote any C -rough geodesic joining points $x, x' \in X$ by $[x, x']$, even though it may not be unique.

Let (X, d) be a metric space and $x, x', y \in X$. The Gromov product $(x|x')_y$ measures the failure of the triangle inequality to be an equality.

$$(2.2) \quad (x|x')_y = \frac{1}{2} (d(x, y) + d(y, x') - d(x, x')).$$

Definition 2.1. Let $\delta \geq 0$. A roughly geodesic metric space X is δ -hyperbolic if for all $x, x', x'', y \in X$ one has

$$(2.3) \quad (x|x'')_y \geq \min\{(x|x')_y, (x'|x'')_y\} - \delta.$$

X is called *Gromov hyperbolic* or just *hyperbolic* if it is δ -hyperbolic for some δ .

In the context of geodesic metric spaces, one uses the following as the definition of hyperbolicity.

Definition 2.2. A geodesic space X has δ -thin triangles if for all points $x, y, z \in X$ and any choice of geodesics $[x, y], [y, z], [x, z]$, each side of the triangle is in the δ -neighborhood of the other two sides:

$$(2.4) \quad [x, z] \subset N_\delta([x, y] \cup [y, z]).$$

X is said to have thin triangles if its triangles are δ -thin for some δ .

Lemma 2.3 (see [2]). *Let X be a geodesic space. Then X has thin triangles if and only if it is hyperbolic.*

Bonk-Schramm linked the study of roughly geodesic hyperbolic spaces to that of geodesic hyperbolic spaces with the following theorem.

Theorem 2.4 (Bonk-Schramm [5]). *Let X be a roughly geodesic δ -hyperbolic space. Then X embeds isometrically into a geodesic δ -hyperbolic space.*

Corollary 2.5. *Let X be a C -roughly geodesic δ -hyperbolic space. There exists $\delta' > 0$ so that all C -roughly geodesic triangles of X are δ' -thin.*

Proof. It is easy to show that there exists a $D > 0$ depending only on C and δ so that in every δ -hyperbolic geodesic metric space, any C -rough geodesic is at most distance D from a geodesic with the same endpoints.

Consider an embedding $X \hookrightarrow Y$ of X into a geodesic δ -hyperbolic space. A C -roughly geodesic triangle in X can then be approximated by a geodesic triangle in Y with the same vertices. By Lemma 2.3, the approximating triangle is thin, so the original C -roughly geodesic triangle must also be thin. \square

2.2. Boundaries of Hyperbolic Spaces.

Definition 2.6. Let X be a roughly geodesic hyperbolic space. The *boundary at infinity* $\partial_\infty X$ is defined as the set of all roughly geodesic rays in X , considered equivalent if they are at a bounded Hausdorff distance from each other.

Remark 2.7. There are three additional standard ways of defining the boundary at infinity:

- (1) Geodesic rays in X , equivalent if they are at bounded Hausdorff distance from each other.
- (2) Quasi-geodesic rays in X , with the same equivalence.

- (3) Sequences of points in X whose Gromov product converges to ∞ . Two sequences are equivalent if their Gromov product converges to ∞ .

For equivalence of the three definitions in geodesic spaces, see [2]. Equivalence with our definition follows from Section 5 of [5].

Definition 2.8. Let X be a roughly geodesic hyperbolic space, $p \in X$, and $a > 0$. The (a, p) -visual quasimetric on $\partial_\infty X$ is defined by

$$(2.5) \quad d_{p,a}([\xi], [\zeta]) = \lim_{t \rightarrow \infty} a^{-(\xi(t)|\zeta(t))_p}.$$

Any metric bi-Lipschitz to $d_{p,a}$ for some p and a is called an (a) -visual metric.

It is well-known that the limit in (2.5) exists. There exists an $a_0 > 1$ so that for any $a \geq a_0$ the function $d_{p,a}$ is bi-Lipschitz to a metric (see [10] Chapter 7). For CAT(-1) spaces, $d_{p,a}$ is itself a metric. For different basepoints p , the corresponding functions $d_{p,a}$ are bi-Lipschitz equivalent, so that the family of visual metrics does not depend on p .

Definition 2.9. Let X be a roughly geodesic hyperbolic space, η a roughly geodesic ray, and $a > 0$. The $([\eta], a)$ -parabolic visual quasimetric is defined by

$$(2.6) \quad d_{[\eta],a}([\xi], [\zeta]) = \lim_{t \rightarrow \infty} a^{t - (\xi(t), \zeta(t))_{\eta(t)}}.$$

Any metric on $\partial_\infty X \setminus \{[\eta]\}$ bi-Lipschitz to $d_{[\eta],a}$ for some a is called an (a) -parabolic visual metric.

As for visual metrics, there exists an a_0 so that for any $a > a_0$ the function $d_{[\eta],a}$ is bi-Lipschitz to a metric (see [10] Chapter 8).

Remark 2.10. In general, there is no way to identify punctured boundaries corresponding to different punctures $[\eta] \in \partial_\infty X$. While visual metrics based inside X have bounded diameter (since the Gromov product is non-negative), parabolic visual metrics are unbounded, with the puncture $[\eta]$ positioned infinitely far from the other points. Shanmugalingam-Xie related parabolic visual metrics to visual metrics by means of metric inversions in [20].

Because the Gromov product is bounded below by 0, visual metrics are bounded. Bonk-Schramm show that every complete bounded metric space arises as the visual boundary of some hyperbolic space (see Section 2.4). We are primarily interested in the case of manifolds with large isometry groups. In this case, one has more control on the boundary.

Heintze showed in [13] that every homogeneous manifold M with negative curvature is diffeomorphic to $N \rtimes \mathbb{R}^+$ with a left-invariant metric, where N is a nilpotent Lie group. There is a left-invariant metric d_N on N so that in the semi-direct product structure \mathbb{R}^+ acts on N by homotheties. It is known (cf. Theorem 3.19) that $\partial_\infty M$ can be identified with the one-point compactification of N and that under this identification d_N is a parabolic visual metric. As a nilpotent Lie group, N is diffeomorphic to \mathbb{R}^n for some n , so $\partial_\infty M$ is a sphere.

Of particular interest is the case when M is a non-compact rank one symmetric space. For real hyperbolic space $\mathbb{H}_{\mathbb{R}}^{n+1}$, the round metric on S^n serves as a visual metric on $\partial_\infty \mathbb{H}_{\mathbb{R}}^{n+1}$. The Euclidean metric on \mathbb{R}^n is a parabolic visual metric on the

punctured boundary of $\mathbb{H}_{\mathbb{R}}^{n+1}$. For complex hyperbolic space $\mathbb{H}_{\mathbb{C}}^{n+1}$, these metrics are no longer visual. Instead, one gives S^{2n+1} the sub-Riemannian metric induced by the inclusion $S^{2n+1} \hookrightarrow \mathbb{C}^{2n+2}$ to obtain a visual metric on $\partial_{\infty}\mathbb{H}_{\mathbb{C}}^{n+1}$. Likewise, \mathbb{R}^{2n+1} is given the structure of the Heisenberg group and the corresponding Carnot-Carathéodory metric to obtain a parabolic visual metric on the punctured boundary. See Section 4.2 for more details.

2.3. Quasi-Isometries and Quasi-Symmetries.

Definition 2.11. Let (X, d_X) and (Y, d_Y) be metric spaces, $C \geq 0$, and $L \geq 1$. An (L, C) -quasi-isometry from X to Y is a map $f : X \rightarrow Y$ such that for all $x, y \in X$

$$(2.7) \quad -C + L^{-1}d_X(x, y) \leq d_Y(fx, fy) \leq Ld_X(x, y) + C$$

and furthermore Y is covered by an L -neighborhood of $f(X)$. If the latter condition is not satisfied, we say f is a quasi-isometry *into* Y . Note that in both cases f need not be continuous, injective or surjective.

Definition 2.12. Let $f : X \rightarrow Y$ be a quasi-isometry between δ -hyperbolic metric spaces. If Y is geodesic and proper, then the Arzela-Ascoli theorem implies that the image of every geodesic ray under a quasi-isometry is uniformly close to a geodesic ray. One may then define the induced map $\partial_{\infty}f$ on the boundary as $\partial_{\infty}f([\gamma]) := [\gamma']$, where γ' is some geodesic ray close to $f(\gamma)$. If Y is not proper, the Arzela-Ascoli Theorem does not apply, and the construction of $\partial_{\infty}f$ is more involved; see [5].

Definition 2.13. Let (X, d_X) and (Y, d_Y) be metric spaces and η an increasing bijection $\eta : [0, \infty) \rightarrow [0, \infty)$. A surjective map $g : X \rightarrow Y$ is an η -quasi-symmetry if for all $x, y, z \in X$ we have

$$(2.8) \quad \frac{d(gx, gy)}{d(gx, gz)} \leq \eta \left(\frac{d(x, y)}{d(x, z)} \right).$$

Remark 2.14. For X and Y equal to Euclidean space or the Heisenberg group, quasi-symmetry is equivalent to quasi-conformality.

Note that for a hyperbolic space X , any two visual metrics on $\partial_{\infty}X$ are quasi-symmetric to each other. However, the class of visual metrics is not closed under quasi-symmetries.

Definition 2.15. Let X be a hyperbolic space. A metric d on $\partial_{\infty}X$ is in the *conformal gauge* if there exists a visual metric d_{vis} on $\partial_{\infty}X$ such that the identity map $id : (\partial_{\infty}X, d) \rightarrow (\partial_{\infty}X, d_{\text{vis}})$ is a quasi-symmetry. The conformal gauge on the punctured boundary $\partial_{\infty}X \setminus \{[\xi]\}$ is, likewise, the class of metrics quasi-symmetrically equivalent to a parabolic visual metric. Note that visual metrics are not quasi-symmetric to parabolic visual metrics, since quasi-symmetries preserve boundedness. They may, however, be quasi-conformally equivalent.

The following is a classical result for geodesic spaces (see [16, 6]). For roughly geodesic spaces, see Theorem 6.5 of [5].

Theorem 2.16. *Let X be a hyperbolic space and $f : X \rightarrow X$ a quasi-isometry. Then $\partial_{\infty}f$ is a quasi-symmetry with respect to any metric in the conformal gauge of $\partial_{\infty}X$. If some point $[\xi] \in \partial_{\infty}X$ is fixed by $\partial_{\infty}f$, then $\partial_{\infty}f$ is a quasi-symmetry with respect to any metric in the conformal gauge of $\partial_{\infty}X \setminus \{[\xi]\}$.*

Sketch of proof. One shows that $\partial_\infty f$ is a quasi-symmetry (in the appropriate sense) for the quasimetrics $d_{p,a}$ and $d_{[\xi],a}$. It is then a quasi-symmetry with respect to any metric in the conformal gauge. \square

For visual metrics, Bonk-Schramm provided a partial converse to Theorem 2.16 in [5] by showing that every quasi-symmetry of $\partial_\infty X$ is induced by a quasi-isometry of X itself. Our main theorem states that under further assumptions on X , every quasi-symmetry of $\partial_\infty X \setminus [\xi]$ is in fact induced by a bi-Lipschitz map on X .

2.4. Recovering the interior. Does every metric space arise as the boundary of some hyperbolic space?

Definition 2.17. Let (X, d_X) be a metric space with diameter $0 < D \leq \infty$. Define a new metric space $\text{Con}(X)$ with underlying set $\text{Con}(X) = X \times (0, D]$ (or $(0, \infty)$ if $D = \infty$) and metric:

$$(2.9) \quad d_{\text{Con}(X)}((x, t), (x', t')) = 2 \log \left(\frac{d_X(x, x') + \max\{t, t'\}}{\sqrt{tt'}} \right).$$

The following theorem combines several results in Sections 7 and 8 of [5]:

Theorem 2.18 (Bonk-Schramm [5]). *Suppose (X, d_X) is a complete metric space of diameter $D \leq \infty$. Then $\text{Con}(X)$ is hyperbolic and roughly geodesic. If $D < \infty$, then under a natural identification of X with $\partial_\infty \text{Con}(X)$, the metric d_X is e -visual.*

We now suppose that X is unbounded and generalize the identification in Theorem 2.18. The following are geodesic rays in $\text{Con}(X)$:

$$(2.10) \quad \gamma_x(s) = (x, e^{-s}) \quad \gamma_{\infty, x_0}(s) = (x_0, e^s),$$

where $s \in [0, \infty)$, and $x, x_0 \in X$. For different choices of x_0 , the rays γ_{∞, x_0} are bounded distance from each other. When the specific x_0 is understood or irrelevant, we will simply write γ_∞ . Given a roughly geodesic ray γ , denote its equivalence class in the boundary by $[\gamma] \in \partial_\infty X$. We will refer to $[\gamma_\infty]$ as the *vertical direction*.

Consider now $\text{Con}(X) = X \times \mathbb{R}^+$ as a subset of $(X \times \mathbb{R}) \cup \{\infty\}$. Identifying a geodesic γ_x with its endpoint on the boundary gives the mapping

$$(2.11) \quad \iota : \overline{X} = X \cup \{\infty\} \rightarrow \partial_\infty \text{Con}(X)$$

$$(2.12) \quad x \mapsto [\gamma_x].$$

It is clear that ι is well-defined and injective. We next show that for some X it is a bijection.

Proposition 2.19. *Let (X, d_X) be an unbounded complete metric space. Suppose there is a homothety α dilating X by a factor $a > 1$ so that for some compact $K \subset X$, $\cup_{n \in \mathbb{Z}} \alpha^n K = X$. Then ι is a bijection.*

Proof. The metric d_X restricts to a metric on K , which we continue to denote d_X . We also have a natural isometric embedding $\text{Con}(K) \hookrightarrow \text{Con}(X)$. Abusing notation, extend α to an isometry of $\text{Con}(X)$ by taking $\alpha(x, t) = (\alpha(x), at)$. Then, $\text{Con}(\alpha^n K) \subset \text{Con}(X)$ for each n , isometric to $\text{Con}(K)$ by α^n . Each $(\alpha^n K, d_X)$ is complete and bounded. By Theorem 2.18 and Corollary 2.5, there exist $C, C' \geq C$ so that $\text{Con}(X)$ is C -rough geodesic and C -rough triangles in $\text{Con}(X)$ are C' -thin.

Suppose that ι is not surjective. Then there exists a rough geodesic ray ξ that is not equivalent to γ_x for any $x \in \overline{X}$. Increasing C and C' if necessary, we may assume that ξ is C -rough and has $\xi(0) = (x_0, 1)$ for some $x_0 \in X$. Write $\xi(s) = (x(s), t(s))$ for $x(s) \in X$, $t(s) \in \mathbb{R}^+$.

We first show that $t(s)$ is bounded above. Suppose $t(s)$ is unbounded. Let $\gamma_\infty = \gamma_{\infty, x_0}$. Since $[\xi] \neq [\gamma_\infty]$, there exists $s_0 > 0$ so that $d(\gamma_\infty(s_0), \xi) > C'$. If $t(s)$ is unbounded, we can select s_1 so that

$$(2.13) \quad \log(t(s_1)) > s_0 + C + C'.$$

Set $s_2 = \log t(s_1)$ so that we have $\gamma_\infty(s_2) = (x_0, t(s_1))$. Fix a C -rough geodesic $[\xi(s_1), \gamma_\infty(s_2)]$.

Consider now the triangle with vertices $(x_0, 1)$, $\xi(s_1)$, and $\gamma_\infty(s_2)$ and edges ξ , γ_∞ , and $[\xi(s_1), \gamma_\infty(s_2)]$. Since this triangle is C' -thin, by choice of s_0 we have

$$(2.14) \quad d(\gamma_\infty(s_0), [\xi(s_1), \gamma_\infty(s_2)]) < C'.$$

We compute some distances using (2.9):

$$\begin{aligned} d(\xi(s_1), \gamma_\infty(s_2)) &= 2 \log \left(\frac{d(x(s_1), x_0) + t(s_1)}{t(s_1)} \right) \\ d(\xi(s_1), \gamma_\infty(s_0)) &= 2 \log \left(\frac{d(x(s_1), x_0) + t(s_1)}{\sqrt{t(s_1)} e^{s_0}} \right) \\ &= 2 \log \left(\frac{d(x(s_1), x_0) + t(s_1)}{t(s_1)} \right) + \log(t(s_1)/e^{s_0}) \end{aligned}$$

Applying (2.13), we conclude

$$d(\xi(s_1), \gamma_\infty(s_0)) > d(\xi(s_1), \gamma_\infty(s_2)) + C + C'.$$

Thus, the C -rough geodesic between $\xi(s_1), \gamma_\infty(s_2)$ cannot come within C' of the point $\gamma_\infty(s_0)$, contradicting 2.14.

We thus have that $t(s)$ is bounded above. Now, by Theorem 2.18, we know ξ is not contained in $\alpha^n \text{Con}(K)$ for any n . Thus, the rough geodesic ray ξ must travel arbitrarily far along X while maintaining a bounded t coordinate. Such a path is also inefficient, as can be shown with a similar calculation as above. Thus, ξ cannot exist and ι is a bijection. □

Corollary 2.20. *Let the notation be as in Proposition 2.19 and identify \overline{X} with $\partial_\infty \text{Con}(X)$ by means of ι . Assume further that α has a fixed point $x_0 \in X$ and that x_0 is an interior point of K . Then the metric d_X is a parabolic visual metric on $\partial_\infty \text{Con}(X) \setminus \{[\gamma_\infty]\}$.*

Proof. Under ι , $K \subset X$ is identified with $\partial_\infty \text{Con}(K) \subset \partial_\infty \text{Con}(X)$. By Theorem 2.18, the metric d_X restricted to $K \subset X$ is a visual metric with base e . We may take $o = (x_0, 1)$ as the basepoint for the Gromov products in (2.5).

Let $x, y \in X$. For large n we have $\alpha^{-n}x, \alpha^{-n}y \in K$ since the distances to $x_0 \in K$ can be made arbitrarily small. We compute for large $n \in \mathbb{N}$, denoting bi-Lipschitz

equivalence independent of n by \approx .

$$\begin{aligned}
d_X(x, y) &= a^n d_X(\alpha^{-n}x, \alpha^{-n}y) && (\alpha \text{ is an } a\text{-homothety}) \\
&\approx \lim_{s \rightarrow \infty} a^n e^{-(\alpha^{-n}\gamma_x(s)|\alpha^{-n}\gamma_y(s))_o} && (\text{definition of visual metric}) \\
&= \lim_{s \rightarrow \infty} a^n e^{-(\gamma_x(s)|\gamma_y(s))_{\alpha^n o}} && (\text{applying the isometry } \alpha^n) \\
&= \lim_{s \rightarrow \infty} e^{n \log a - (\gamma_x(s)|\gamma_y(s))_{\alpha^n o}} && (\text{change of base})
\end{aligned}$$

Note now that $\alpha^n o = \gamma_\infty(n \log a)$. We thus have

$$d_X(x, y) \approx \lim_{s \rightarrow \infty} e^{n \log a - (\gamma_x(s)|\gamma_y(s))_{\gamma_\infty(n \log a)}}.$$

We may take $n = \lfloor s \rfloor / \log a$, so that

$$d_X(x, y) \approx \lim_{s \rightarrow \infty} e^{s - (\gamma_x(s)|\gamma_y(s))_{\gamma_\infty(s)}}.$$

Comparing to (2.6), d_X is a parabolic visual metric.

□

3. METRIC SIMILARITY SPACES

3.1. Motivating Example: Real Hyperbolic Space. We now sketch a proof of the theorem of Tukia and Väisälä that motivated this paper.

Theorem 3.1 ([24]). *Let $f : \mathbb{R}^n \rightarrow \mathbb{R}^n$ be quasi-conformal. Then there exists a quasi-conformal $F : \mathbb{R}^{n+1} \rightarrow \mathbb{R}^{n+1}$ such that the restriction to $\mathbb{R}^n \times \{0\}$ agrees with f . Furthermore, F restricted to the upper half-space is bi-Lipschitz with respect to the hyperbolic metric.*

Proof sketch. Quasi-conformal maps on \mathbb{R}^n are quasi-symmetric. Fix η so that f is η -quasi-symmetric. Denote $H^+ = \mathbb{R}^n \times \mathbb{R}^+$, with the hyperbolic metric (see Section 4.2). Let $K = [0, 1]^n$, $Q_0 = K \times [1, 2]$. Then the elements of $\mathcal{Q} = \{2^a(Q_0 + b) : a \in \mathbb{Z}, b \in \mathbb{Z}^n\}$ are the *tiles* of the dyadic decomposition of H^+ . With respect to the hyperbolic metric, the tiles are isometric to each other by maps that preserve the local combinatorics of the tiling.

Given a homeomorphism $g : \mathbb{R}^n \rightarrow \mathbb{R}^n$, define the *basic extension* $\tilde{g} : H^+ \rightarrow H^+$ by

$$\tilde{g}(x, t) = \left(g(x), \max_{|x-y| \leq t} |g(x) - g(y)| \right).$$

The extension \tilde{g} is a homeomorphism that extends continuously to ∂H^+ , with values on ∂H^+ given by g . The extension \tilde{f} of f is not necessarily bi-Lipschitz on small scales, but this can be repaired using a combination of piecewise-linear (PL) approximation and Sullivan's extension theorem, as we now describe.

Recall that, except in dimension 4, any homeomorphism of Euclidean space may be approximated by a PL homeomorphism. Because H^+ is Riemannian, it is locally bi-Lipschitz to Euclidean space, and Euclidean methods such as PL approximation may be applied locally.

We would like to apply PL approximation to obtain a bi-Lipschitz approximation F of f on each tile $Q \in \mathcal{Q}$. However, the approximations on adjacent tiles need

not agree. To avoid this, let $\mathcal{Q} = \mathcal{Q}_1 \sqcup \dots \sqcup \mathcal{Q}_N$ be a coloring of \mathcal{Q} . On tiles in \mathcal{Q}_1 , define F using a PL approximation of \tilde{f} .

Remark 3.2. In the case $n = 4$, one can first use PL approximation on \mathbb{R}^n and then lift to H^+ while preserving the bi-Lipschitz property. Note that there do exist homeomorphisms of \mathbb{R}^4 to itself that cannot be piecewise-linearly approximated.

The next step is to interpolate the approximations. This is done using Sullivan's bi-Lipschitz extension result (see Lemma 3.36). For each tile in \mathcal{Q}_2 , Sullivan extension provides values of F that agree with the values given on the tiles \mathcal{Q}_1 . Once the interpolation is complete on \mathcal{Q}_1 , continue inductively on the remaining colors until F is defined on all of H^+ .

Every step above used infinitely many bi-Lipschitz approximations. To guarantee that the bi-Lipschitz constant stays bounded, fix $p_0 \in H^+$ and let

$$\widetilde{\mathcal{F}}_\eta = \{\tilde{g} : H^+ \rightarrow H^+ : \tilde{g}(p_0) \in Q_0 \text{ and } g \text{ is } \eta\text{-quasi-symmetric}\}.$$

It follows from the Arzela-Ascoli Theorem that $\widetilde{\mathcal{F}}_\eta$ is compact, so finitely many bi-Lipschitz approximations are sufficient to approximate any map in $\widetilde{\mathcal{F}}_\eta$ on Q_0 . Given another tile Q , there are isometries g, h of H^+ with $gQ = Q_0$, and the composition gh is in $\widetilde{\mathcal{F}}_\eta$.

Thus, in the first step, the PL approximation theorem is invoked only finitely many times. Furthermore, this produces only finitely many configurations (up to isometry) to interpolate in the second step. Thus, for each of the finitely many colors Sullivan's extension theorem is invoked only finitely many times.

The resulting map F is then defined using finitely many bi-Lipschitz maps, and is therefore bi-Lipschitz and furthermore quasi-conformal. Since it is a bounded distance from \tilde{f} , F has the prescribed extension to the boundary. Furthermore, the hyperbolic metric on H^+ is conformally equivalent to the Euclidean metric, so F is a quasi-conformal map with respect to the Euclidean metric. Extending to all of \mathbb{R}^{n+1} via reflection completes the proof. \square

Our goal now is to generalize Theorem 3.1 as Theorem 3.40, applicable to a wider class of spaces, which we next define and study.

3.2. Metric Similarity Spaces: Definition and Structure. We first define stacked tilings. A metric space H homeomorphic to \mathbb{R}^n is said to have a *stacked tiling* if it possesses both a discrete co-compact action by some group Γ and a compatible homothety α . More precisely:

Definition 3.3. Suppose (H, d_H) is a complete metric space homeomorphic to \mathbb{R}^n for some n so that:

- (1) There exists a discrete group $\Gamma \subset \text{Isom}(H)$ acting properly on H and a compact connected domain $K \subset H$ such that $H = \cup\{\gamma K : \gamma \in \Gamma\}$.
- (2) For all $\gamma \in \Gamma$ not equal to the identity, the interiors of K and γK are disjoint.
- (3) There exists a homothety $\alpha : H \rightarrow H$ that rescales distances by some $a > 1$.
- (4) For $\gamma \in \Gamma$, $\alpha\gamma\alpha^{-1} \in \Gamma$.

- (5) There is a finite subset $\Gamma' \subset \Gamma$ with $\alpha K = \cup\{\gamma' K : \gamma' \in \Gamma'\}$.

We then say that H possesses a *stacked tiling*, following the terminology of [21].

Example 3.4. Consider $H = \mathbb{R}^n$ with group $\Gamma = \mathbb{Z}^n$. We may then take K to be the unit cube and $\alpha(x) = 2x$. Then, $2K$ is a cube with side length 2, which is the union of 2^n translates of K .

Problem 3.5. Is Condition 5 equivalent to the subgroup $\alpha\Gamma\alpha^{-1}$ having finite index in Γ ?

Definition 3.6. Suppose H possesses a stacked tiling and furthermore $H^+ = H \times \mathbb{R}^+$ has a Riemannian metric d_{H^+} so that

- (1) For $\gamma \in \Gamma$, the mapping $(x, t) \mapsto (\gamma x, t)$ of H^+ is an isometry.
- (2) The mapping $(x, t) \mapsto (\alpha x, at)$ of H^+ is an isometry.

Under these conditions, we refer to H^+ as an $(n+1)$ -dimensional metric similarity space with base H .

Example 3.7. Let H be as in Example 3.4 and $H^+ = \mathbb{R}^n \times \mathbb{R}^+$. The hyperbolic metric on H^+ is then invariant under both the “parabolic” transformations $(x, t) \mapsto (x + x', t)$ and the “hyperbolic” transformation $(x, t) \mapsto (2x, 2t)$.

Definition 3.8. Denote by Γ_α the semigroup

$$(3.1) \quad \Gamma_\alpha = \{\alpha^n \gamma : n \in \mathbb{Z}, \gamma \in \Gamma\}.$$

Depending on context, we view elements of Γ_α as mappings of either H or H^+ , based on Definition 3.6. Condition 4 of Definition 3.3 provides a relation on the semigroup.

If α commutes with elements of Γ , then Γ_α is a group. Even when this is the case, Γ_α is not discrete. However, the stacked tiling of H can be turned into a tiling of H^+ by disjoint tiles related by Γ_α . More specifically:

Definition 3.9. Let $Q_0 = K \times [1, a]$. The *tiles* of H^+ are the compact sets

$$(3.2) \quad \mathcal{Q} = \{gQ_0 : g \in \Gamma_\alpha\}.$$

We refer to Q_0 as the *fundamental tile* of H^+ . For each $n \in \mathbb{Z}$, $\{\alpha^n \gamma Q_0 : \gamma \in \Gamma\}$ is the n^{th} layer of \mathcal{Q} . The tiles are isometric, and we have $\cup \mathcal{Q} = H^+$. Furthermore, the interiors of any two tiles are disjoint.

A *coloring* of \mathcal{Q} is a partition $\mathcal{Q} = \mathcal{Q}_1 \sqcup \cdots \sqcup \mathcal{Q}_N$ such that two tiles of the same color (index) do not intersect. Because the tiles \mathcal{Q} are in bijective correspondence with Γ_α , every coloring of the tiles induces a coloring $\Gamma_1 \sqcup \cdots \sqcup \Gamma_N$ of Γ_α .

Remark 3.10. The fundamental tile Q_0 serves for a metric similarity space the role played by a fundamental domain in a space with a co-compact group action. However, even if Γ_α is a group, it is not discrete and Q_0 is not a fundamental domain.

In general, elements of Γ_α do not take tiles to tiles.

Definition 3.11. Let $Q \in \mathcal{Q}$. We call an element $\gamma \in \Gamma_\alpha$ a *combinatorial isometry with respect to Q* if for any tile Q' with $Q \cap Q' \neq \emptyset$, the image $\gamma Q'$ of Q' is a tile. We will not mention Q explicitly if it is obvious which tile is relevant.

Lemma 3.12. *Let $\gamma \in \Gamma$ and $Q \in \mathcal{Q}$. If Q is in the n^{th} layer of \mathcal{Q} , for $n \leq 0$, then $\gamma Q \in \mathcal{Q}$ is a tile.*

Proof. A tile in the n^{th} layer has the form $\alpha^n \gamma Q_0$. If $n \leq 0$, we can find using Condition 4 of Definition 3.3 $\gamma_1 \in \Gamma$ so that $\gamma \alpha^n \gamma Q_0 = \alpha^n \gamma_1 Q_0$, again a tile. \square

Lemma 3.13 (Normalization Lemma). *Let $Q \in \mathcal{Q}$. Then $Q = g\gamma'Q_0$ for some $g \in \Gamma_\alpha$ and $\gamma' \in \Gamma'$, with g a combinatorial isometry.*

Proof. Because α takes tiles to tiles, we may assume Q is in the 0^{th} layer of the tiling, so that $Q = \gamma_0 K \times [1, a]$ for some $\gamma_0 \in \Gamma$.

We now consider the stacked tiling of H . Because the tiles $\{\gamma K : \gamma \in \Gamma\}$ tile H , so do the tiles $\{\alpha\gamma K : \gamma \in \Gamma\}$. There then exists $\alpha\gamma_1 K$ with $\text{int}(\alpha\gamma_1 K) \cap \text{int}(\gamma_0 K) \neq \emptyset$, where int denotes the interior of a tile. Taking $g = \alpha\gamma_1\alpha^{-1} \in \Gamma$, we have:

$$\alpha\gamma_1 K = g\alpha K = g \left(\bigcup_{\gamma' \in \Gamma'} \gamma' K \right).$$

Since $\text{int}(\alpha\gamma_1 K) \cap \text{int}(\gamma_0 K) \neq \emptyset$, we have $\gamma_0 = g\gamma'$.

Returning to H^+ , we can write $Q = g\gamma'Q_0$. Tiles intersecting Q are in either the same layer or one of the adjacent layers. By Lemma 3.12, $g = \alpha\gamma_1\alpha^{-1}$ takes tiles each of these three layers to other tiles in the same layer. \square

Corollary 3.14. *There exists a coloring of \mathcal{Q} .*

Proof. Every $Q \in \mathcal{Q}$ is equivalent to $\gamma'Q_0$ for some $\gamma' \in \Gamma'$ by a combinatorial isometry. There are finitely many tiles $\gamma'Q_0$, each with finite valence. Thus, the valence of the adjacency graph is bounded, and it is therefore clear that it has a coloring. \square

Corollary 3.15. *Let $\Gamma_\alpha = \Gamma_1 \sqcup \cdots \sqcup \Gamma_N$ be a coloring induced from a coloring of \mathcal{Q} . Then there exists an open neighborhood Q_1 of Q_0 such that for any i and distinct $g, h \in \Gamma_i$, $gQ_1 \cap hQ_1 = \emptyset$.*

Proof. Let $\gamma' \in \Gamma'$. Let $D_{\gamma'} = d_{H^+}(\gamma'Q_0, (\gamma'Q_0)^*)$, where

$$(\gamma'Q_0)^* = \cup\{Q \in \mathcal{Q} : Q \cap \gamma'Q_0 \neq \emptyset\}.$$

Setting $D = \min_{\gamma' \in \Gamma'} D_{\gamma'} \neq 0$, let Q_1 be the $D/2$ -neighborhood of Q_0 . By Lemma 3.13, we have considered all the tiles up to combinatorial isometry, so Q_1 has the desired property. \square

3.3. Base and Boundary. We next define a metric on \mathcal{Q} using its adjacency graph and use a variation of the Milnor-Švarc Lemma ([7] p. 140) to find a quasi-isometry between H^+ and $\text{Con}(H)$. This allows us to conclude that H^+ is hyperbolic and identify $H \cup \{\infty\}$ with $\partial_\infty H^+$.

Definition 3.16. Define a graph with vertices \mathcal{Q} and an edge between two vertices $Q \neq Q'$ if $Q \cap Q' \neq \emptyset$. By an abuse of notation, we continue to refer to this graph as \mathcal{Q} . Denote by $d_{\mathcal{Q}}$ the graph metric on \mathcal{Q} which assigns to each edge length 1.

Lemma 3.17. *Let $H^+ = H \times \mathbb{R}^+$ be a metric similarity space and $\text{Con}(H)$ as in Definition 2.17. Then $\text{id} : H^+ \rightarrow \text{Con}(H)$ is a quasi-isometry.*

Proof. Recall that $d_{\text{Con}(H)}$ is ϵ -roughly geodesic for some $\epsilon > 0$. The metric d_{H^+} is complete and Riemannian, so geodesic and in particular ϵ -roughly geodesic. Let d be either one of these metrics on the set $H \times \mathbb{R}^+ = H^+ = \text{Con}(H)$. Note that Γ_α acts on (H^+, d) with isometries for either choice of d . Increasing the fundamental tile Q_0 and taking the corresponding sub-semi-group of Γ_α if necessary, we may assume that $\epsilon \ll D = 3 \text{diam}(Q_0)$.

Define a function $f : H^+ \rightarrow \mathcal{Q}$ as follows. For each $x \in H^+$, let $f(x) \in \mathcal{Q}$ be a tile containing x . We now show that f is a quasi-isometry between (H^+, d) and the graph \mathcal{Q} . From this, we will be able to conclude that $\text{id} : (H^+, d_{H^+}) \rightarrow (\text{Con}(H), d_{\text{Con}(H)})$ is a quasi-isometry.

By definition of D , we know that if $x, y \in X$ and $d(f(x), f(y)) = 1$, then $d(x, y) \leq 2D/3$. Conversely, it is easy to see that there is an $N > 0$ so that if $d(x, y) < D$, then $d_{\mathcal{Q}}(f(x), f(y)) \leq N$.

Let $x, y \in H^+$ and let $n = d_{\mathcal{Q}}(f(x), f(y))$. Pick a geodesic $f(x) = Q_1, \dots, Q_{n+1} = f(y)$ in \mathcal{Q} , and choose points q_1, \dots, q_{n+1} in H^+ with $f(q_i) = Q_i$. We then have $d(q_i, q_{i+1}) \leq 2D/3$ so by the triangle inequality we have

$$d(x, y) \leq 2nD/3 = 2d_{\mathcal{Q}}(f(x), f(y))D/3.$$

Thus, f is co-Lipschitz.

Suppose now $x, y \in H^+$ are joined by an ϵ -rough geodesic $r(s)$, with $s \in [0, d(x, y)]$. Let $r(s_0), \dots, r(s_n)$ be a minimal sequence of points such that $s_0 = 0$, $s_n = d(x, y)$, and $d(r(s_i), r(s_{i+1})) \leq D$. We estimate

$$\begin{aligned} d_{\mathcal{Q}}(f(x), f(y)) &\leq d_{\mathcal{Q}}(f(r(s_0)), f(r(s_1))) + \dots + d_{\mathcal{Q}}(f(r(s_{n-1})), f(r(s_n))) \\ &\leq Nn \leq N \left(\frac{d(x, y)}{D - \epsilon} + 1 \right), \end{aligned}$$

where the last inequality follows from the fact that r is an ϵ -rough geodesic and the choice of s_i is minimal.

Since f is onto, \mathcal{Q} is quasi-isometric to both d_{H^+} and $d_{\text{Con}(H)}$. Thus, the identity map from H^+ to $\text{Con}(H)$ is a quasi-isometry. \square

Lemma 3.18. *Let $f : X \rightarrow Y$ be a quasi-isometry between hyperbolic spaces. Suppose d is in the conformal gauge of $\partial_\infty X$. Then the induced metric $(\partial_\infty f)_* d$ on ∂Y is in the conformal gauge of ∂Y . The same holds for parabolic conformal gauges.*

Proof. Since d is in the conformal gauge of $\partial_\infty X$, it is bi-Lipschitz to a visual metric d_1 . Let d_2 be a visual metric on $\partial_\infty Y$. By Theorem 2.16, the induced metric $(\partial_\infty f)_* d_1$ on $\partial_\infty Y$ is quasi-symmetric to the metric d_2 . Thus, same holds for $(\partial_\infty f)_* d$. \square

The following follows directly from Lemma 3.17 and Lemma 3.18:

Theorem 3.19. *Let H^+ be a metric similarity space with base H . Then H^+ is Gromov hyperbolic and the mapping $\iota : H \cup \{\infty\} \rightarrow \partial_\infty H^+$ defined in (2.11) is a bijection. Under this identification, the metric d_H is in the parabolic conformal gauge of $\partial_\infty H^+ \setminus \{[\gamma_\infty]\}$.*

Corollary 3.20. *Let $H^+ = H \times (0, \infty)$ be a metric similarity space. Under the identification $H \times \{0\} = \partial_\infty H^+ \setminus \{\infty\}$, every quasi-isometry of H^+ to itself extends continuously to a map on $H \times [0, \infty)$.*

Remark 3.21. In view of Theorem 3.19 and Corollary 3.20, we use the notation ∂H^+ to denote both $\partial_\infty H^+$ and $H \cup \{\infty\}$. Likewise, for a quasi-isometry f of H^+ , ∂f denotes the extension of f to ∂H^+ .

3.4. Extension from the Base.

Definition 3.22. Let H^+ be an n -dimensional metric similarity space with base H . A map $\wedge : C(H, H) \rightarrow C(H^+, H^+)$, denoted by $f \mapsto \widehat{f}$, is a *lifting method* if

- (1) If f is quasi-symmetric, then \widehat{f} extends to a homeomorphism $\widehat{f} : H \times [0, \infty) \rightarrow H \times [0, \infty)$, with $\widehat{f}|_{H \times \{0\}} = f$.
- (2) \wedge is continuous with respect to the compact-open topologies on the two spaces,
- (3) \wedge is Γ_α -equivariant: $\widehat{g_1 f g_2} = g_1 \widehat{f} g_2$ for all $g_1, g_2 \in \Gamma_\alpha$.

Lemma 3.23. *Let F be a quasi-isometry of H^+ preserving the vertical direction. If $\widehat{\partial F}$ is a quasi-isometry, then it is bounded distance from F .*

Proof. Quasi-isometries are injective on the large scale, so there is a quasi-isometry F^{-1} so that $F^{-1} \circ F$ is bounded distance from the identity map. Since $\partial(\widehat{\partial F}) = \partial F$, the boundary map $\partial(F^{-1} \circ \widehat{\partial F})$ is the identity map. It then follows from the stability of quasi-geodesics that $F^{-1} \circ \widehat{\partial F}$ is finite distance from the identity map. \square

We now discuss a specific lifting method \sim introduced by Tukia-Väisälä. Others are also available, as mentioned in the introduction.

Definition 3.24. Given $f : H \rightarrow H$ quasi-symmetric, let $\widetilde{f} : H^+ \rightarrow H^+$ defined by

$$(3.3) \quad \widetilde{f}(x, t) = (f(x), \tau(x, t)) \quad \tau(x, t) = \max_{|x-y| \leq t} |f(x) - f(y)|.$$

The following is easy to see from the definition of τ :

Lemma 3.25. *The mapping $\sim : f \mapsto \widetilde{f}$ is a lifting method.*

We next give a sequence of examples demonstrating properties of \sim .

Example 3.26. Let $f : \mathbb{R}^n \rightarrow \mathbb{R}^n$ be given by $f(x) = \lambda x$ for $\lambda \in \mathbb{R}$. Then $\widetilde{f}(x, t) = (\lambda x, \lambda t)$, a hyperbolic transformation of the upper half-space $\mathbb{H}_{\mathbb{R}}^{n+1}$. More generally, if $F : \mathbb{H}_{\mathbb{R}}^{n+1} \rightarrow \mathbb{H}_{\mathbb{R}}^{n+1}$ is an isometry preserving the point at infinity, then $\widetilde{\partial F} = F$.

Example 3.27. Let $f : \mathbb{R}^2 \rightarrow \mathbb{R}^2$ be given by $f(x, y) = (x, 2y)$. Then $\tilde{f}(x, y, t) = (x, 2y, 2t)$ is 2-bi-Lipschitz, as can be seen by applying \tilde{f} to a geodesic and computing the effect on its speed. By extension, any affine map $f : \mathbb{R}^n \rightarrow \mathbb{R}^n$ lifts to a $|\Lambda/\lambda|$ -bi-Lipschitz map \tilde{f} , where Λ is the eigenvalue of maximum norm and λ is the eigenvalue of minimum norm.

Example 3.28. Let $f : \mathbb{R} \rightarrow \mathbb{R}$ be given by $f(x) = x^3$, a quasi-symmetric map. Then $\tilde{f} : \mathbb{H}_{\mathbb{R}}^2 \rightarrow \mathbb{H}_{\mathbb{R}}^2$ is given by $\tilde{f}(x, t) = (x^3, (|x| + t)^3 - |x|^3)$. Away from $x = 0$, \tilde{f} is differentiable, and the differential becomes singular as x approaches 0. Since the metric on $\mathbb{H}_{\mathbb{R}}^2$ is Riemannian, this implies \tilde{f} is not bi-Lipschitz.

Example 3.28 shows that a quasi-symmetry f does not necessarily lift to a bi-Lipschitz map under \sim . Theorem 1.3 implies, however, that for every quasi-symmetry f , the lift \tilde{f} is a quasi-isometry and may in fact be approximated by a bi-Lipschitz map.

3.5. Compactness of Quasi-Symmetries. Let $H^+ = H \times \mathbb{R}^+$ be a metric similarity space, and the notation as in Section 3.2. Fix a homeomorphism $\eta : [0, \infty) \rightarrow [0, \infty)$, distinct points $x_0, y_0 \in H$, and a lifting method Λ .

Definition 3.29. A mapping $f : H \rightarrow H$ is *normalized* if $f(x_0) \in K$ and

$$d_H(x_0, y_0) \leq d_H(fx_0, fy_0) \leq ad_H(x_0, y_0)$$

for a, K as in Definition 3.3 for H . We also define:

$$\begin{aligned} \mathcal{F}_\eta &= \{f : H \rightarrow H : f \text{ is a normalized } \eta\text{-quasi-symmetry}\} \subset C(H, H), \\ \widehat{\mathcal{F}}_\eta &= \{\widehat{f} : f \in \mathcal{F}_\eta\} \subset C(H^+, H^+). \end{aligned}$$

Note that $\widehat{\mathcal{F}}_\eta$ is the image of \mathcal{F}_η under the continuous lifting map Λ .

Remark 3.30. For any η -quasi-symmetric map $f : H \rightarrow H$, there exist $n \in \mathbb{Z}$ and $\gamma \in \Gamma$ so that $(\gamma\alpha^n f) \in \mathcal{F}_\eta$ and $(\gamma\alpha^n \widehat{f}) \in \widehat{\mathcal{F}}_\eta$. Note that $\gamma\alpha^n$ need not be in Γ_α if $n > 0$.

Lemma 3.31. \mathcal{F}_η and $\widehat{\mathcal{F}}_\eta$ are compact in the compact-open topology.

Proof. It suffices to prove that \mathcal{F}_η is compact, because $\widehat{\mathcal{F}}_\eta$ is its image under a continuous map. By the Arzela-Ascoli theorem for metric spaces, [12] Theorem 10.28, it suffices to show that \mathcal{F}_η is closed, equicontinuous and has compact orbits. We switch to the notation $|x - y| := d_H(x, y)$ and follow the discussion on p. 85 of [12].

The family \mathcal{F}_η is defined by closed conditions, except possibly the condition that the maps be η -quasi-symmetric embeddings. It follows from the Arzela-Ascoli theorem that every limit of η -quasi-symmetric maps is a η -quasi-symmetric embedding or constant. The condition $|x_0 - y_0| \leq |fx_0 - fy_0|$ rules out the possibility of a limit map that is constant. Thus \mathcal{F}_η is closed.

For any $z \in H$ and $f \in \mathcal{F}_\eta$, we have by definition of quasi-symmetry:

$$\frac{|fz - fx_0|}{|fy_0 - fx_0|} \leq \eta\left(\frac{|z - x_0|}{|y_0 - x_0|}\right),$$

so that

$$|fz - fx_0| \leq \eta \left(\frac{|z - x_0|}{|y_0 - x_0|} \right) |fy_0 - fx_0| \leq a |x_0 - y_0| \eta \left(\frac{|z - x_0|}{|y_0 - x_0|} \right).$$

Since $f(x_0) \in K$, $|fz - fx_0|$ is bounded for each $z \in H$. Since H is complete and homeomorphic to \mathbb{R}^n , this implies that the closure of $\mathcal{F}_\eta(z)$ is compact for every $z \in H$.

We now prove equicontinuity at every $z_0 \neq x_0 \in H$. Let $z \in H$ and $f \in \mathcal{F}_\eta$.

$$\begin{aligned} |fz - fz_0| &\leq \eta \left(\frac{|z - z_0|}{|x_0 - z_0|} \right) |fx_0 - fz_0| \\ &\leq \eta \left(\frac{|z - z_0|}{|x_0 - z_0|} \right) \eta \left(\frac{|z - z_0|}{|x_0 - y_0|} \right) |x_0 - y_0| \\ &\leq \eta \left(\frac{|z - z_0|}{|x_0 - z_0|} \right) \eta \left(\frac{|x_0 - z_0|}{|x_0 - y_0|} \right) |fx_0 - fy_0| \\ &\leq a |x_0 - y_0| \eta \left(\frac{|z - z_0|}{|x_0 - z_0|} \right) \eta \left(\frac{|x_0 - z_0|}{|x_0 - y_0|} \right). \end{aligned}$$

Thus, $|fz - fz_0|$ is controlled by a continuous function of $|z - z_0|$. For $z_0 = x_0$, the same calculation holds with x_0 and y_0 switched. Thus, \mathcal{F}_η is equicontinuous. \square

Corollary 3.32. *Up to specified additive error, $\widehat{\mathcal{F}}_\eta$ contains finitely many elements.*

3.6. Bi-Lipschitz approximation. We first re-phrase a standard result about Riemannian manifolds.

Lemma 3.33. *Let H^+ be a metric similarity space of dimension $n + 1$. Then H^+ is locally bi-Lipschitz to \mathbb{R}^{n+1} .*

Proof. Recall that H^+ is isometric to \mathbb{R}^{n+1} with a Riemannian metric d_{H^+} . We claim that on any compact $K \subset \mathbb{R}^{n+1}$, d_{H^+} is bi-Lipschitz to the Euclidean metric d_E .

Let $D = \max\{\text{diam}_{H^+}(K), \text{diam}_E(K)\}$, and let $K' = N_{H^+}(K, D) \cup N_E(K, D)$ denote the union of the closed D -neighborhoods of K with respect to each metric. Note that K' contains the union of all geodesic segments from K to itself with respect to both metrics. Since both metrics are complete, K' is compact. Next, denote by $T^1(K') \simeq K' \times S^{n-1}$ the (compact) unit tangent bundle to K' as a subset of \mathbb{R}^n , with respect to either norm. Denote by ds_E and ds_{H^+} the line elements for the corresponding metrics and set

$$L = \max \left\{ \sup_{v \in T^1(K')} \frac{ds_{H^+}(v)}{ds_E(v)}, \sup_{v \in T^1(K')} \frac{ds_E(v)}{ds_{H^+}(v)} \right\}.$$

By definition of d_E and d_{H^+} as Riemannian metrics, they are L -bi-Lipschitz equivalent. \square

The following follows immediately from standard piecewise-linear approximation results (see [19] p. 194) and Lemma 3.33.

Lemma 3.34. *Let H^+ be a metric similarity space of dimension $n \neq 4$, $f : H^+ \rightarrow H^+$ a homeomorphism, and $K \subset H^+$ compact. Then for any $\epsilon > 0$ there exists a locally bi-Lipschitz $g : H^+ \rightarrow H^+$ such that $|g - f|_K < \epsilon$.*

Remark 3.35. There exist counterexamples to Lemma 3.34 in dimension 4.

Lemma 3.34 provides only local approximations of an embedding. Stitching local approximations together into a global bi-Lipschitz embedding is a difficult task. The only tool that is currently available for this purpose is the deep extension theorem of Sullivan. See [22] for the original statement and [23] for a more detailed exposition of the proof. We state the Tukia-Väisälä adaptation of Sullivan's result in the context of metric similarity spaces.

Lemma 3.36 ([22, 24]). *Let H^+ be a metric similarity space. Suppose $U \subset H^+$ has compact closure and \mathcal{F} is a compact family of embeddings of U into H^+ that have locally bi-Lipschitz approximations. Let $V \subset U$, $U' \subset U$ open with $\overline{U'} \subset U$, $V' \subset V$ with $\overline{V'} \subset V$. For any $\epsilon > 0$, there exists $\delta > 0$ such that:*

If $f \in \mathcal{F}$ and $g : V \rightarrow H^+$ is a bi-Lipschitz embedding with $d_{H^+}(f, g) < \delta$, then there is a bi-Lipschitz embedding $g' : U \rightarrow H^+$ that extends the approximation of f in the sense that:

- (1) $g' = g$ on V' ,
- (2) $d_{H^+}(g', f) < \epsilon$ on U' .

Remark 3.37. To paraphrase, if g is an approximation of f on V , then the approximation can be extended to U , at the cost of taking subsets of U and V . The resulting error is at most ϵ , as long as the original error was less than δ . Note that δ depends on the compact family \mathcal{F} in which f is contained, but not on f itself.

Remark 3.38. For metric similarity spaces of dimension not equal to 4, Lemma 3.36 applies to the families $\widehat{\mathcal{F}}_\eta$, by Lemma 3.31 and Lemma 3.34.

Once we have approximated an injective map by a locally bi-Lipschitz map, we would like for it to remain an embedding. The following theorem guarantees that for every function sufficiently close to an element of $\widehat{\mathcal{F}}_\eta$, local injectivity implies global injectivity.

Lemma 3.39. *Let H^+ be a metric similarity space, $r > 0$, and $\widehat{\mathcal{F}}_\eta$ a family of normalized mappings of H^+ . Then there exists $\epsilon > 0$ satisfying the following condition:*

Let $h : H^+ \rightarrow H^+$ with $|h - \widehat{f}| < \epsilon$ for some $\widehat{f} \in \widehat{\mathcal{F}}_\eta$. Suppose that for any $z, w \in H^+$ with $|z - w| < r$ we have $h(z) \neq h(w)$. Then h is injective.

Proof. Recall that Q_0 denotes the compact fundamental tile of H^+ and that the family $\widehat{\mathcal{F}}_\eta$ is compact by Lemma 3.31. Since H^+ is a complete Riemannian space, its spheres $S(z, r) = \{w \in H^+ : d_{H^+}(z, w) \leq r\}$ are compact. We claim the lemma holds for

$$\epsilon = \frac{1}{3} \min \left\{ d_{H^+}(\widehat{f}z, \widehat{f}w) : \widehat{f} \in \widehat{\mathcal{F}}_\eta, z \in Q_0, w \in S(z, r) \right\} > 0.$$

Indeed, let $d_{H^+}(h, \hat{f}) < \epsilon$, and $z, w \in H^+$. If $d_{H^+}(z, w) < r$, we are done, so we assume $d_{H^+}(z, w) \geq r$. If $z \in Q_0$, then for topological reasons we have

$$d_{H^+}(\hat{f}z, \hat{f}w) \geq d_{H^+}(\hat{f}z, \hat{f} \cdot S(z, r)) \geq 3\epsilon.$$

Since $d_{H^+}(\hat{f}z, hz) < \epsilon$ and $d_{H^+}(\hat{f}w, hw) < \epsilon$, we have $d_{H^+}(h(z), h(w)) \geq \epsilon$.

If $z \notin Q_0$, then $z \in g_1 Q_0$ for some $g_1 \in \Gamma_\alpha$. By Remark 3.30, there is also an isometry g_2 of H^+ so that $g_2 f g_1^{-1} \in \widehat{\mathcal{F}}_\eta$. The calculation in the previous paragraph then applies, with $g_2 f g_1^{-1}$ substituted for \hat{f} , $g_1 z$ for z , and $g_1 w$ for w .

We then have that for all $z, q \in H^+$ with $d_{H^+}(z, w) \geq r$, $d_{H^+}(hz, hw) > \epsilon$. In particular, h is injective. \square

3.7. Main Theorem. We now prove the main theorem. For an overview of the proof in the case $H^+ = \mathbb{H}_{\mathbb{R}}^{n+1}$, see Section 3.1.

Theorem 3.40 (Main Theorem). *Let $H^+ = H \times \mathbb{R}^+$ be a metric similarity space of dimension not equal to 4, \wedge a lifting method, $\epsilon > 0$, and $f : H \rightarrow H$ a quasi-symmetry. There exists a bi-Lipschitz embedding $F : H^+ \rightarrow H^+$ with $d_{H^+}(\hat{f}, F) < \epsilon$, and $\partial F = f$.*

Proof. Recall from Section 3.2 that $H^+ = \cup \mathcal{Q}$, a union of tiles with disjoint interiors. The tiles \mathcal{Q} are translates by the semigroup Γ_α of the fundamental tile Q_0 . By Lemma 3.14, the tiling \mathcal{Q} of H^+ is equipped with a coloring $\mathcal{Q} = \mathcal{Q}_1 \sqcup \cdots \sqcup \mathcal{Q}_N$, inducing a coloring $\Gamma_\alpha = \Gamma_1 \sqcup \cdots \sqcup \Gamma_N$.

We will approximate \hat{f} one tile at a time, first by piecewise-linear approximation on \mathcal{Q}_1 , and using bi-Lipschitz extension on the other colors $i = 2, \dots, N$. While there are infinitely many tiles of each color, we will use Lemma 3.13 to argue that PL approximation and bi-Lipschitz extension only need to be invoked finitely many times per color.

Step 1 (Base case: PL approximation). We will first approximate \hat{f} on \mathcal{Q}_1 within error ϵ_1 , to be specified later. By Corollary 3.15, there exists a neighborhood Q_1 of Q_0 with compact closure whose translates $\Gamma_1 Q_1$ are disjoint. We now apply PL approximation (Lemma 3.34) on the disjoint domains $\Gamma_1 Q_1$.

Let $\gamma_1 \in \Gamma_1$. By Remark 3.30, there exists $n \in \mathbb{N}$, $\gamma \in \Gamma$ so that $h \circ \hat{f} \circ \gamma_1^{-1} \in \widehat{\mathcal{F}}_\eta$, where $h\gamma\alpha^n$ is an isometry of H^+ . The family $\widehat{\mathcal{F}}_\eta$ is compact, so by Corollary 3.32 contains only finitely many maps (up to error $\epsilon/2$). Let f' be one of these maps, equivalent to $h \circ \hat{f} \circ \gamma_1^{-1}$ up to an error of $\epsilon/2$. Approximating f' on Q_1 within error $\epsilon/2$ using Lemma 3.34, we get a bi-Lipschitz approximation F_{γ_1} of $h \circ \hat{f} \circ \gamma_1^{-1}$ with error at most ϵ_1 . Returning to the original tile, $h^{-1} \circ F_{\gamma_1} \circ \gamma_1$ approximates \hat{f} on the tile $\gamma_1 Q_1$.

We repeat in the same way on the remaining domains $\Gamma_1 Q_1$. The first approximation F_1 of \hat{f} is now defined on $\cup \Gamma_1 Q_1$, and approximates \hat{f} with error at most ϵ_1 .

Step 2 (Inductive step: Sullivan extension). Let $i > 1$ and suppose \hat{f} has been approximated by a map F_{i-1} on $\cup_{j=1}^{i-1} \Gamma_j Q_{i-1}$, where Q_{i-1} is a neighborhood of Q_0

with compact closure. Assume furthermore that F_{i-1} is defined by finitely many maps. More precisely, assume there is a finite collection of bi-Lipschitz maps \mathcal{F}_{i-1} such that if $\gamma \in \Gamma_1 \cup \dots \cup \Gamma_{i-1}$ and h is an isometry of H^+ such that $h \circ \hat{f} \circ \gamma^{-1} \in h \circ \hat{f} \circ \gamma^{-1}$, then $h \circ F_{i-1} \circ \gamma^{-1} \in \mathcal{F}_{i-1}$.

Let Q_i be a neighborhood of Q_0 in H^+ with compact closure with $\overline{Q_i} \subset Q_{i-1}$. By the choice of Q_1 , the domains $\Gamma_i Q_i$ are disjoint.

Let $\gamma \in \Gamma_i$. The map \hat{f} is defined on all of γQ_{i-1} . The previous approximation F_{i-1} is defined on $\gamma Q_{i-1} \cap (\cup_{j=1}^{i-1} \Gamma_j Q_{i-1})$. Applying Lemma 3.36 with

$$\begin{aligned} U &= \gamma Q_{i-1}, & V &= \gamma Q_{i-1} \cap (\cup_{j=1}^{i-1} \Gamma_j Q_{i-1}) \\ U' &= \gamma Q_i, & V' &= \gamma Q_i \cap (\cup_{j=1}^{i-1} \Gamma_j Q_i) \end{aligned}$$

and $\epsilon = \epsilon_i$ we get a bi-Lipschitz embedding F_γ defined on γQ_i and compatible with F_{i-1} on $\cup_{j=1}^{i-1} \Gamma_j Q_i$. Note that Lemma 3.36 applies because of Remark 3.30 and Lemma 3.34. Repeat for all $\gamma \in \Gamma_i$ to get a new approximation F_i of \hat{f} on $\cup_{j=1}^i \Gamma_j Q_j$.

By the Finiteness Lemma 3.13 and the inductive finiteness assumption, we need apply Lemma 3.36 only finitely many times. Thus, the finiteness assumption is preserved.

Step 3 (Completing the Induction). By induction, $F = F_N$ is defined on all of H^+ . Up to isometries of H^+ , it is piecewise defined by elements of a finite family \mathcal{F}_N of bi-Lipschitz maps. Thus, it is uniformly locally bi-Lipschitz. Note that a locally bi-Lipschitz map need not be globally injective.

Step 4 (Approximation Error and Injectivity). There are two demands on the approximation error. The first is the maximal error ϵ given by the statement of the theorem. Secondly, we would like for F to be injective. Let $r = d_{H^+}(\partial Q_0, \partial Q_N) > 0$. If $x, y \in H^+$ and $d(x, y) < r$, then x and y are both contained in an enlarged tile γQ_N for some $\gamma \in \Gamma_\alpha$. Thus, we know the approximation F is injective on the scale r . Lemma 3.39 provides an ϵ_N such that F is globally injective if $d_H(F, \hat{f}) < \epsilon_N$. We may assume $\epsilon_N < \epsilon$.

Inductively, suppose ϵ_i is already defined. Lemma 3.36, when applied in the i^{th} step of the inductive argument above, provides an approximation of error at most ϵ_i as long as the previous error is at most $\epsilon_{i-1} < \epsilon_i$, for some ϵ_i . We require this degree of error in the previous step.

The PL approximation base case automatically complies with error requirement ϵ_1 .

With these error bounds on the approximation process, F is injective. An injective uniformly locally bi-Lipschitz mapping of a path metric space is globally bi-Lipschitz.

Step 5 (Boundary). The maps F and \hat{f} differ by bounded additive noise. It then follows from Theorem 3.19 that $\partial F = f$. \square

Corollary 3.41. *Let H^+ be a metric similarity space of dimension $n \neq 4$ with base H and $f : H \rightarrow H$ a quasi-symmetry. Then for any lifting method \wedge , the mapping $\hat{f} : H^+ \rightarrow H^+$ is a quasi-isometry.*

Corollary 3.42. *Let H^+ be a metric similarity space of dimension $n \neq 4$. Then any quasi-isometry f of H^+ preserving the vertical direction is bounded distance from a bi-Lipschitz map. If the action of $\text{Isom}(H^+)$ on ∂H^+ is transitive, then f need not preserve the vertical direction.*

Proof. By Theorem 3.19, ∂f is a quasi-symmetry of H . Applying Theorem 3.40 with any lifting method (e.g. the lifting method \sim) gives a bi-Lipschitz mapping F of H^+ with $\partial F = \partial f$. It follows from the stability of quasi-geodesics that the distance between f and F is bounded.

If $\text{Isom}(H^+)$ acts on ∂H^+ transitively, then one may compose with the appropriate isometry in order to preserve the vertical direction. \square

4. DISCUSSION

In this section, we first present some examples of metric similarity spaces and then the connection to past results and open problems.

4.1. Homogeneous Negatively Curved Spaces. Let M be a homogeneous negatively curved manifold. Heintze showed in [13] that M is diffeomorphic to a Lie group $N \rtimes_{\alpha_s} \mathbb{R}^+$, where N is some nilpotent group modeled on \mathbb{R}^n and \mathbb{R}^+ acts on \mathbb{R}^n by linear transformations α_s . The group N possesses a left-invariant metric for which α_s is a homothety of factor s . We first provide an example for $N = \mathbb{R}^n$. See Sections 4.2 and 4.3 for additional examples.

Example 4.1. Let $\mathbb{R}^n \rtimes \mathbb{R}^+$ be a Lie group modeled on $\mathbb{R}^n \times \mathbb{R}^+$ with group structure

$$(4.1) \quad (x, t) \cdot (x', t') = (x + e^{At} x', tt'),$$

where A is a diagonal matrix with positive eigenvalues $\lambda_1, \dots, \lambda_n$. Like any Lie group, $\mathbb{R}^n \rtimes_A \mathbb{R}^+$ has a natural left-invariant metric. Define a metric d_A on \mathbb{R}^n by

$$(4.2) \quad |(x_1, \dots, x_n)|_A = |x_1|^{1/\lambda_1} + \dots + |x_n|^{1/\lambda_n} \quad d_A(x, x') = |x - x'|_A.$$

Proposition 4.2. *Suppose A is a diagonal matrix with positive integer eigenvalues $\lambda_1, \dots, \lambda_n$. Then the Lie group $\mathbb{R}^n \rtimes_A \mathbb{R}^+$ with any left-invariant Riemannian metric is a metric similarity space with base (\mathbb{R}^n, d_A) ,*

Proof. Let $\Gamma = \mathbb{Z}^n$ and $K = [0, 1]^n$. The mapping

$$\alpha(x_1, \dots, x_n) = (2^{\lambda_1} x_1, \dots, 2^{\lambda_n} x_n)$$

of \mathbb{R}^n dilates distances by a factor of 2. The image of K under α is

$$(4.3) \quad \alpha K = [0, 2^{\lambda_1}] \times \dots \times [0, 2^{\lambda_n}],$$

tilled by finitely many copies of K . Furthermore, α commutes with Γ . Thus, (\mathbb{R}^n, d_A) has a stacked tiling. The metric on $\mathbb{R}^n \rtimes_A \mathbb{R}^+$ is by construction invariant under the transformations Γ_α , so $\mathbb{R}^n \rtimes_A \mathbb{R}^+$ is a metric similarity space. \square

Definition 4.3. Let G be a Lie group modeled on $\mathbb{R}^{n_1} \times \dots \times \mathbb{R}^{n_r}$, with elements $x \in G$ written as $x = x_{ij}$ with $1 \leq i \leq r, 1 \leq j \leq n_i$, and group structure of the form

$$(4.4) \quad (x \cdot x')_{ij} = x_{ij} + x'_{ij} + F_{ij}(x, x').$$

The group G is a *rational graded nilpotent group of step r* if each F_{ij} is a polynomial in the variables $x_1, \dots, x_{i-1}, x'_1, \dots, x'_{i-1}$ with rational coefficients and satisfies

$$(4.5) \quad F_{ij}(\delta_s x, \delta_s x') = s^i F_{ij}(x, x'),$$

where the transformation δ_s is given by

$$(4.6) \quad (\delta_s x)_{ij} = s^i x_{ij}.$$

The *Carnot-Carathéodory* metric on N is defined as follows. Let \mathcal{D} be the left-invariant distribution that at the origin corresponds to the \mathbb{R}^{n_1} subspace: $\mathcal{D}_0 = \mathbb{R}^{n_1}$. Call a path $\gamma(t)$ *horizontal* if the velocity $\gamma'(t) \in \mathcal{D}_{\gamma(t)}$ for almost all t . Give \mathcal{D} a left-invariant inner product $g_{\mathcal{D}}$. The Carnot-Carathéodory metric d_{CC} between two points is defined as the $g_{\mathcal{D}}$ -length of the shortest horizontal curve between the points. By Chow's Theorem, d_{CC} is well-defined. It is clear that δ_s is a homothety of d_{CC} by factor s .

Example 4.4. Both \mathbb{R}^n and the Heisenberg group \mathcal{H}^n (see Section 4.2) are rational graded nilpotent groups. See [21] for further examples.

Theorem 4.5 (Strichartz [21]). *Every rational graded nilpotent group possesses a stacked tiling.*

Corollary 4.6. *Let N be a rational graded nilpotent group. Let \mathbb{R}^+ act on N by the dilations δ_s in Definition 4.3. Then $N \rtimes_{\delta_s} \mathbb{R}^+$ with any left-invariant Riemannian metric is a metric similarity space.*

Proof. The group $N \rtimes_{\delta_s} \mathbb{R}^+$ acts on itself by isometries. By 4.5, if one picks the homothety $\alpha = \delta_2$ of N , there is a corresponding subgroup $\Gamma \subset N$ and a compact $K \subset N$ satisfying the definition of stacked tiling. \square

Remark 4.7. Every homogeneous negatively curved manifold decomposes as $N \rtimes_{\alpha_s} \mathbb{R}^+$ for some family of diffeomorphisms α_s . However, even if N is a rational graded nilpotent group, generally $\alpha_s \neq \delta_s$. See Example 4.1.

4.2. Rank One Symmetric Spaces. Among the homogeneous negatively curved manifolds, the most-studied ones are the non-compact rank one symmetric spaces. These are the real, complex, and quaternionic hyperbolic spaces, as well as the octonionic plane. In this section, we show that both real hyperbolic space $\mathbb{H}_{\mathbb{R}}^{n+1}$ and complex hyperbolic space $\mathbb{H}_{\mathbb{C}}^{n+1}$ are metric similarity spaces. The same is true of the quaternionic and octonionic spaces. However, according to Pansu rigidity [17], any quasi-isometry of the quaternionic and Cayley spaces is in fact an isometry, and any quasi-symmetry of the boundary is conformal. There exist quasi-isometries of $\mathbb{H}_{\mathbb{R}}^{n+1}$ and $\mathbb{H}_{\mathbb{C}}^{n+1}$ that are not equivalent to isometries (see [15]).

Definition 4.8. Real hyperbolic space $\mathbb{H}_{\mathbb{R}}^{n+1}$ is the space $\mathbb{R}^n \rtimes_A \mathbb{R}^+$ defined in Example 4.1, where A is taken to be the identity matrix. The line element is explicitly given by

$$ds^2 = \frac{dx^2 + dt^2}{t^2}.$$

Proposition 4.9. *Real hyperbolic space $\mathbb{H}_{\mathbb{R}}^{n+1}$ is a metric similarity space with base (\mathbb{R}^n, d_E) , where d_E is the Euclidean metric.*

Proof. This follows directly from Proposition 4.2. \square

We now describe complex hyperbolic space $\mathbb{H}_{\mathbb{C}}^n$. The definition corresponds to that of the Klein disk model of $\mathbb{H}_{\mathbb{R}}^n$, but uses complex coefficients instead of real ones. For more details, see the book of Goldman [11] and a recent survey of Parker-Platis [18].

Denote by \mathbb{CP}^n complex projective space of dimension n . Let J be a Hermitian form on \mathbb{C}^{n+1} of type $(n, 1)$, so $J(z, z) \in \mathbb{R}$ for all z . Complex hyperbolic space is the set of negative points in \mathbb{CP}^n with respect to J (we will define a metric shortly).

$$\mathbb{H}_{\mathbb{C}, J}^n = \{p \in \mathbb{CP}^n : J(z, z) < 0 \text{ for } z \in \mathbb{C}^{n+1} \text{ representing } p\}.$$

Because J is unique up to a change of basis, the space is well-defined up to a choice of coordinates, and we write $\mathbb{H}_{\mathbb{C}}^n$ if J is not specified. Taking $z = (z_0, \dots, z_n)$, we define two Hermitian forms:

$$\begin{aligned} J_1(z) &= -|z_0|^2 + |z_1|^2 + \dots + |z_n|^2, \\ J_2(z) &= z_0 \overline{z_1} + z_1 \overline{z_0} + |z_2|^2 + \dots + |z_n|^2. \end{aligned}$$

In the coordinate patch $z_0 = 1$, $\mathbb{H}_{\mathbb{C}, J_1}^n$ is the unit ball in \mathbb{C}^n , while $\mathbb{H}_{\mathbb{C}, J_2}^n$ is the Siegel (or paraboloid) model of complex hyperbolic space:

$$\mathbb{H}_{\mathbb{C}, J_2}^n = \{(z_1, \dots, z_n) : -2\operatorname{Re}(z_1) > |z_2|^2 + \dots + |z_n|^2\}.$$

Let $z, w \in \mathbb{C}^{n+1}$ represent two points $p, q \in \mathbb{H}_{\mathbb{C}, J}^n$. The metric on $\mathbb{H}_{\mathbb{C}, J}^n$ is defined by

$$\cosh^2 d_{\mathbb{H}_{\mathbb{C}, J}^n}(p, q) = \frac{J(z, w)J(w, z)}{J(z, z)J(w, w)}.$$

This metric is, in fact, Riemannian. It may alternately be defined by taking the standard inner product at the origin of the ball model $\mathbb{H}_{\mathbb{C}, J_1}^n$ and uniquely extending to a metric tensor on $\mathbb{H}_{\mathbb{C}, J_1}^n$ invariant under the group of projective transformations preserving J_1 .

Remark 4.10. Unlike real hyperbolic space, complex hyperbolic space does not have constant curvature: for $n > 1$ the sectional curvature κ of $\mathbb{H}_{\mathbb{C}}^n$ varies in the range $-4 \leq \kappa \leq -1$. In the case $n = 1$, $\mathbb{H}_{\mathbb{C}, J_1}^1$ is precisely the Poincaré model of the real hyperbolic plane, rescaled to have constant curvature -4 .

Let \mathcal{H}^{n-1} denote the 2-step nilpotent Heisenberg group (see also [8, 14]) modeled on $\mathbb{C}^{n-1} \times \mathbb{R}$ with group law

$$(4.7) \quad (\zeta, u) \cdot (\zeta', u') = (\zeta + \zeta', u + u' + 2\operatorname{Im}\langle \zeta, \zeta' \rangle),$$

where $\langle \zeta, \zeta' \rangle$ is the standard Hermitian pairing of complex vectors. The Heisenberg group is a rational graded nilpotent group of step 2. As such, it has a Carnot-Carathéodory metric with dilations δ_s (see Definition 4.3).

Define the horospherical model of complex hyperbolic space as $\mathbb{H}_{\mathbb{C}, \text{horo}}^n = \mathcal{H}^{n-1} \times \mathbb{R}^+ = \mathbb{C}^{n-1} \times \mathbb{R} \times \mathbb{R}^+$. Give $\mathbb{H}_{\mathbb{C}, \text{horo}}^n$ a Riemannian metric by identifying it with the Siegel model by means of the diffeomorphism

$$f : \mathbb{H}_{\mathbb{C}, \text{horo}}^n \rightarrow \mathbb{H}_{\mathbb{C}, J_2}^n \subset \mathbb{C} \times \mathbb{C}^{n-1} \quad f : (\zeta, u, v) \mapsto (-v + \mathbf{i}u + |\zeta|^2, \sqrt{2}\zeta).$$

In horospherical coordinates, we have the following relationship between $\mathbb{H}_{\mathbb{C}}^n$ and \mathcal{H}^{n-1} :

- (1) Isometries of \mathcal{H}^{n-1} extend trivially to parabolic isometries of $\mathbb{H}_{\mathbb{C}, \text{horo}}^n$.
- (2) Similarities δ_s of \mathcal{H}^{n-1} extend to hyperbolic isometries defined by $(\zeta, u, v) \mapsto (s\zeta, s^2u, s^2v)$.

Furthermore, we have the following properties of \mathcal{H}^{n-1} .

- (1) The subgroup $\Gamma \subset \mathcal{H}^n$ consisting of elements with integer coordinates is discrete and co-compact.
- (2) By Theorem 4.5 or the results of [3], there exists a “Heisenberg cube”, a fundamental domain for Γ that provides \mathcal{H}^{n-1} with a stacked tiling with $\alpha = \delta_2$.

Proposition 4.11. *Complex hyperbolic space $\mathbb{H}_{\mathbb{C}}^n$ is a metric similarity space with base \mathcal{H}^{n-1} with its Caranot-Carathéodory metric.*

Proof. The horospherical model $\mathbb{H}_{\mathbb{C}}^n$ satisfies the definition, up to taking the square root of the v coordinate, which does not scale properly as stated. \square

4.3. Comparison with a Result of Xie. A Hadamard manifold is a complete negatively curved manifold with pinched sectional curvature. A general Hadamard manifold is not a metric similarity space. However, the following result has significant overlap with our main theorem:

Theorem 4.12 (Xie [25]). *Let M be a Hadamard manifold that is the universal cover of a compact manifold of dimension not equal to 4. Then every quasi-isometry from M to itself is bounded distance from a bi-Lipschitz map of M to itself.*

In particular, Theorem 4.12 applies to complex hyperbolic space and other negatively curved homogeneous manifolds admitting co-compact discrete group actions. We present two examples to distinguish our main Theorem 3.40 from Theorem 4.12.

The following well-known example shows that Theorem 4.12 does not apply to most metric similarity spaces constructed in Example 4.1:

Proposition 4.13. *Let $\mathbb{R}^n \rtimes_A \mathbb{R}^+$ be as in Proposition 4.2, with the λ_i distinct positive integers. Then $\mathbb{R}^n \rtimes_A \mathbb{R}^+$ does not admit a co-compact isometric group action.*

Proof. Consider the group of isometries $\text{Isom}(\mathbb{R}^n \rtimes_A \mathbb{R}^+)$. It is easy to see by considering the Riemannian metric on $\mathbb{R}^n \rtimes_A \mathbb{R}^+$ that in this case the stabilizer $\text{Stab}(p)$ of a point is finite.

Suppose that a discrete group G acts co-compactly on $\mathbb{R}^n \rtimes_A \mathbb{R}^+$. There would then be a finite-index subgroup G' of G whose elements are all either parabolic or hyperbolic. In particular, the left-invariant vector field along the \mathbb{R}^+ direction would descend to $\mathbb{R}^n \rtimes_A \mathbb{R}^+/G'$. The resulting flow on $\mathbb{R}^n \rtimes_A \mathbb{R}^+/G'$ would be isometric but not volume preserving, which is impossible. \square

The next proposition provides a perturbation method for metric similarity spaces.

Proposition 4.14. *Let H^+ be a metric similarity space with base H , a Riemannian metric g , and $Q_0 = K \times [1, a]$. Let f be a positive bump function on the interval*

$[1, a]$, equal to 1 near the endpoints. Extend f to a function on $(0, \infty)$ invariant under the mapping $x \mapsto ax$. Then (H^+, fg) is a metric similarity space.

Proof. If g is complete and invariant under the semigroup Γ_α , then so is fg . \square

Remark 4.15. A more interesting version of Proposition 4.14 would allow arbitrary perturbations of the Riemannian metric within the fundamental tile Q_0 . For the sake of simplicity, we required Γ_α to act on H^+ by isometries when we defined metric similarity spaces. However, our results remain valid if $g \in \Gamma_\alpha$ is only required to be an isometry on some fixed neighborhood of Q_0 . Thus, one may introduce arbitrary bumps in the metric, including the appearance of positive curvature.

4.4. Concluding Remarks. Theorem 3.40 states that every quasi-symmetry of the base of a metric similarity space of dimension not equal to 4 extends to a bi-Lipschitz map of the space itself. This leaves open the following problem.

Problem 4.16. Let H^+ be a 4-dimensional metric similarity space with base H . Is it true that every quasi-symmetry of H^+ preserving the vertical direction is bounded distance from a bi-Lipschitz map? Equivalently, does every quasi-symmetry of H lift to a bi-Lipschitz map of H^+ ?

The exclusion of 4-dimensional spaces from the main theorem is due to the dimension restriction in Lemma 3.34. Because there do exist homeomorphisms of \mathbb{R}^4 that cannot be approximated piecewise-linearly (see [19]), the restriction cannot be removed. However, such pathological homeomorphisms are difficult to construct. This suggests the following problem.

Problem 4.17. Let H^+ be a 4-dimensional metric similarity space with base H . Let \wedge be a lifting method, $\epsilon > 0$, and $f : H \rightarrow H$ a quasi-symmetric embedding. Can \hat{f} be approximated by a bi-Lipschitz map within error ϵ ?

In the case of $H^+ = \mathbb{H}_{\mathbb{R}}^4$, one bypasses the dimension restriction of Lemma 3.34 using PL approximation in \mathbb{R}^3 and the following property of the lifting method \sim defined by (3.3).

Lemma 4.18 (Tukia-Väisälä [24]). *Let $f : \mathbb{R}^n \rightarrow \mathbb{R}^n$ be a homeomorphism. Then f is locally bi-Lipschitz in the Euclidean metric if and only if f is locally bi-Lipschitz in the hyperbolic metric on $\mathbb{H}_{\mathbb{R}}^{n+1}$.*

Lemma 4.18 does not generalize directly to the setting of metric similarity spaces, as its proof uses the fact that the hyperbolic metric is locally bi-Lipschitz to the l^1 sum of the Euclidean metrics on \mathbb{R}^n and \mathbb{R}^+ .

Problem 4.19. Let H^+ be a metric similarity space with base H^+ . Generalize Lemma 4.18 for some lifting method \wedge . More specifically, relate the analytic properties of $f : H \rightarrow H$ and those of $\hat{f} : H^+ \rightarrow H^+$.

Shanmugalingam-Xie showed in [20] that for many nilpotent groups N , all quasi-symmetries of N are in fact bi-Lipschitz. While quasi-symmetries of the Heisenberg group need not be bi-Lipschitz, Capogna and Tang constructed a large class of bi-Lipschitz quasi-symmetries of \mathcal{H}^n in [9]. An appropriate analogue of 4.18 would allow these to be lifted to bi-Lipschitz mappings of $\mathbb{H}_{\mathbb{C}}^{n+1}$. The existence of quasi-symmetries of \mathcal{H}^n that are not bi-Lipschitz suggests the following problem.

Problem 4.20. Develop an analogue of piecewise-linear approximation theory for \mathcal{H}^n . Following the Pansu-Rademacher theorem [17], the phrase “piecewise linear” should generalize to “piecewise horizontally homomorphic”, where a horizontal homomorphism of \mathcal{H}^n preserves the standard splitting of the tangent spaces of \mathcal{H}^n .

One may also attempt a more direct approach to proving the main theorem for $H^+ = \mathbb{H}_{\mathbb{C}}^2$. Beurling-Ahlfors originally proved Theorem 3.1 in the case of $H^+ = \mathbb{H}_{\mathbb{R}}^2$ [4] using an explicit integral formula. While Theorem 3.1 is proven using a different approach, perhaps an explicit bi-Lipschitz lifting may be found for $H^+ = \mathbb{H}_{\mathbb{C}}^2$.

Ahlfors’ proof in [1] of Theorem 3.1 in the case of $\mathbb{H}_{\mathbb{R}}^3$ was based on a decomposition of a quasi-conformal mappings. Given a quasi-conformal mapping f of \mathbb{R}^2 , one finds a quasi-conformal flow f_t on \mathbb{R}^2 so that $f = f_1$. One can then write f as the composition of mappings with small distortion, which can then be approximated by bi-Lipschitz maps using results concerning almost-conformal mappings.

Problem 4.21. For $n \geq 3$, can a quasi-conformal mapping of \mathbb{R}^n be embedded in a quasi-conformal flow? For $n \geq 1$, can a quasi-conformal mapping of \mathcal{H}^n be embedded in a quasi-conformal flow? See [15] for a study of quasi-conformal flows on \mathcal{H}^n and their lifts to $\mathbb{H}_{\mathbb{C}}^{n+1}$.

REFERENCES

1. L. Ahlfors, *Extension of quasiconformal mappings from two to three dimensions*, Proc. Nat. Acad. Sci. U.S.A. **51** (1964), 768–771. MR 0167617 (29 #4889)
2. J. M. Alonso, T. Brady, D. Cooper, V. Ferlini, M. Lustig, M. Mihalik, M. Shapiro, and H. Short, *Notes on word hyperbolic groups*, Group theory from a geometrical viewpoint (Trieste, 1990), World Sci. Publ., River Edge, NJ, 1991, Edited by Short, pp. 3–63. MR 1170363 (93g:57001)
3. Z. Balogh, R. Hofer-Isenegger, and J. Tyson, *Lifts of Lipschitz maps and horizontal fractals in the Heisenberg group*, Ergodic Theory Dynam. Systems **26** (2006), no. 3, 621–651. MR 2237461 (2007i:28009)
4. A. Beurling and L. Ahlfors, *The boundary correspondence under quasiconformal mappings*, Acta Math. **96** (1956), 125–142. MR 0086869 (19,258c)
5. M. Bonk and O. Schramm, *Embeddings of Gromov hyperbolic spaces*, Geom. Funct. Anal. **10** (2000), no. 2, 266–306. MR 1771428 (2001g:53077)
6. M. Bourdon, *Structure conforme au bord et flot géodésique d’un CAT(−1)-espace*, Enseign. Math. (2) **41** (1995), no. 1-2, 63–102. MR 1341941 (96f:58120)
7. M. R. Bridson and A. Haefliger, *Metric spaces of non-positive curvature*, Grundlehren der Mathematischen Wissenschaften [Fundamental Principles of Mathematical Sciences], vol. 319, Springer-Verlag, Berlin, 1999. MR 1744486 (2000k:53038)
8. L. Capogna, D. Danielli, S. Pauls, and J. Tyson, *An introduction to the Heisenberg group and the sub-Riemannian isoperimetric problem*, Progress in Mathematics, vol. 259, Birkhäuser Verlag, Basel, 2007. MR 2312336 (2009a:53053)
9. L. Capogna and P. Tang, *Uniform domains and quasiconformal mappings on the Heisenberg group*, Manuscripta Math. **86** (1995), no. 3, 267–281. MR 1323792 (96f:30019)
10. É. Ghys and P. de la Harpe (eds.), *Sur les groupes hyperboliques d’après Mikhael Gromov*, Progress in Mathematics, vol. 83, Birkhäuser Boston Inc., Boston, MA, 1990, Papers from the Swiss Seminar on Hyperbolic Groups held in Bern, 1988. MR 1086648 (92f:53050)
11. W. M. Goldman, *Complex hyperbolic geometry*, Oxford Mathematical Monographs, The Clarendon Press Oxford University Press, New York, 1999, Oxford Science Publications. MR 1695450 (2000g:32029)
12. J. Heinonen, *Lectures on analysis on metric spaces*, Springer, New York, 2000.
13. E. Heintze, *On homogeneous manifolds of negative curvature*, Math. Ann. **211** (1974), 23–34. MR 0353210 (50 #5695)

14. A. Korányi and H. M. Reimann, *Foundations for the theory of quasiconformal mappings on the Heisenberg group*, Adv. Math. **111** (1995), no. 1, 1–87. MR 1317384 (96c:30021)
15. ———, *Equivariant extension of quasiconformal deformations into the complex unit ball*, Indiana Univ. Math. J. **47** (1998), no. 1, 153–176. MR 1631549 (99g:32059)
16. G. Margulis, *The isometry of closed manifolds of constant negative curvature with the same fundamental group.*, Dokl. Akad. Nauk. SSSR **192** (1970), 736–737.
17. P. Pansu, *Métriques de Carnot-Carathéodory et quasiisométries des espaces symétriques de rang un*, Ann. of Math. (2) **129** (1989), no. 1, 1–60. MR 979599 (90e:53058)
18. J. R. Parker and I. D. Platis, *Complex hyperbolic quasi-Fuchsian groups*, Geometry of Riemann surfaces, London Math. Soc. Lecture Note Ser., vol. 368, Cambridge Univ. Press, Cambridge, 2010, pp. 309–355. MR 2665016
19. T. B. Rushing, *Topological embeddings*, Academic Press, New York, 1973, Pure and Applied Mathematics, Vol. 52. MR 0348752 (50 #1247)
20. N. Shanmugalingam and X. Xie, *A rigidity property for some negatively curved solvable groups*, (2011), Preprint.
21. R. S. Strichartz, *Self-similarity on nilpotent Lie groups*, Geometric analysis (Philadelphia, PA, 1991), Contemp. Math., vol. 140, Amer. Math. Soc., Providence, RI, 1992, pp. 123–157. MR 1197594 (94e:43011)
22. D. Sullivan, *Hyperbolic geometry and homeomorphisms*, Geometric topology (Proc. Georgia Topology Conf., Athens, Ga., 1977), Academic Press, New York, 1979, pp. 543–555. MR 537749 (81m:57012)
23. P. Tukia and J. Väisälä, *Lipschitz and quasiconformal approximation and extension*, Ann. Acad. Sci. Fenn. Ser. A I Math. **6** (1981), no. 2, 303–342 (1982). MR 658932 (84a:57016)
24. P. Tukia and J. Väisälä, *Quasi-conformal extension from dimension n to $n + 1$* , Annals of Mathematics **115** (1982), 331–348.
25. X. Xie, *Quasiisometries between negatively curved Hadamard manifolds*, J. Lond. Math. Soc. (2) **79** (2009), no. 1, 15–32. MR 2472131 (2009i:53029)

UNIVERSITY OF ILLINOIS URBANA-CHAMPAIGN, 1407 W. GREEN STREET, URBANA, IL 61802.
 EMAIL: ANTON@LUKYANENKO.NET